

The Uniform Distribution in Incentive Dynamics

Dashiell E.A. Fryer

The University of Illinois at Urbana-Champaign

July 3, 2012

Abstract

The uniform distribution is an important counterexample in game theory as many of the canonical game dynamics have been shown not to converge to the equilibrium in certain cases. In particular none of the canonical game dynamics converge to the uniform distribution in a form of rock-paper-scissors where the amount an agent can lose is more than the agent can win, despite this fact, it is the unique Nash equilibrium. I will show that certain incentive dynamics are asymptotically stable at the uniform distribution when it is an incentive equilibrium.

Incentive dynamics [Fry12] are given by

$$\dot{x}_{i\alpha} = \varphi_{i\alpha}(x) - x_{i\alpha} \sum_{\beta} \varphi_{i\beta}(x)$$

where $\varphi(x)$ is the a valid incentive for the game. It was shown that if the incentive for a finite game is continuous, there exists a fixed point characterized by

$$\varphi_{i\alpha}(\hat{x}) = \hat{x}_{i\alpha} \sum_{\beta} \varphi_{i\beta}(\hat{x}) \quad \forall \alpha, i.$$

Notice that if this occurs at the uniform distribution, either $\varphi_{i\alpha}(\hat{x})$ are all zero, or they are all the same for each agent.

Nash's original incentive function is fixed if and only if all the component incentives are zero and thus it can only be in the first case described above. In contrast, the incentive function given by $\varphi_{i\alpha}^D(x) = \sum_{\gamma} (a_{\alpha\gamma} - u_i(x))_+$ is only zero when $u_i(x) \geq \max_{\gamma} u_i(e_{\alpha}, e_{\gamma})$, where $e_{\gamma} \in S_{-i}$ which can occur at the uniform distribution only if the game is constant, which is a degenerate case of little interest. Despite their differences we will demonstrate that the two can agree under certain circumstances. Also, we will see that the latter incentive is globally asymptotically stable at a uniform Nash equilibrium where the canonical dynamics fail to converge.

0.1 A Bad Game of Rock-Paper-Scissors

The standard game of Rock-Paper-Scissors (RPS) is given as a two person zero sum game with payoffs given in the table below on the left.

0, 0	-1, 1	1, -1	0, 0	-b, a	a, -b
1, -1	0, 0	-1, 1	a, -b	0, 0	-b, a
-1, 1	1, -1	0, 0	-b, a	a, -b	0, 0

To the right of the RPS payoffs we have a generalized RPS with a and b both positive. The case when $b > a$, or an agent can lose more than it can win, is an important example of a game. The unique Nash equilibrium for this game is the uniform distribution. We have seen many examples of incentive dynamics that have Nash equilibrium as their interior fixed points, such as the replicator equations, projection dynamics, the logit equations, best reply dynamics, and the Brown-von Neumann-Nash equations. However, in every

one of these cases the dynamics do not converge to the unique equilibrium as shown in the figures below¹. This leads us to the natural question: does any incentive dynamic converge to a rest point from any initial point?

1 Agreement Among Incentives

We note the incentive $\varphi_{i\alpha}^D(x) = \sum_{\gamma} (a_{\alpha\gamma} - u_i(x))_+$ can be rewritten in the form $\varphi_{i\alpha}^D(x) = \sum_{\gamma} (u_i(e_{\alpha}, x_{-i}) - u_i(x) + a_{\alpha\gamma} - u_i(e_{\alpha}, x_{-i}))_+$, which shows that it is similar to a Nash comparison in that we are checking the payoff given the other agents' strategies are fixed. However, we are tempering that comparison by taking away the amount by which the agent is not receiving a preferred payoff available in the game. We will now show there is a class of games, which includes general RPS, with the property that the uniform distribution is a Nash equilibrium as well as an incentive equilibrium for $\varphi^D(x)$. First we will need the following lemma.

Lemma 1. *If A_i is the payoff matrix for the i th agent, then \hat{x} is a Nash equilibrium where each agent is using the uniform distribution over its strategies if and only if for each i , A_i has an equal sum across rows.*

Proof. We begin by noting that for an interior Nash equilibrium we must have

$$u_i(e_1, x_{-i}) = \dots = u_i(e_{s_i}, x_{-i}), \quad \forall i.$$

It should suffice then to calculate the value of just one of the $u_i(e_{\alpha}, x_{-i})$. We

¹The images were produced using the Dynamo Mathematica package developed by Sandholm, Dokumaci, and Franchetti [SDF11]. Colors indicate speed: blue is slowest and red is fastest

will use the n -linearity of the payoffs to complete the task.

$$u_1(e_1, x_{-1}) = \sum_{j_2=1}^{s_2} \frac{1}{s_2} u_1(e_1, e_{j_2}, x_3, \dots, x_n) \quad (1)$$

$$= \frac{1}{s_2} \sum_{j_2=1}^{s_2} u_1(e_1, e_{j_2}, x_3, \dots, x_n) \quad (2)$$

$$= \frac{1}{s_2} \frac{1}{s_3} \sum_{j_2=1}^{s_2} \sum_{j_3=1}^{s_3} u_1(e_1, e_{j_2}, e_{j_3}, x_4, \dots, x_n) \quad (3)$$

$$= \dots = \frac{1}{\prod_{i \in N \setminus \{1\}} s_i} \sum_{j_2=1}^{s_2} \sum_{j_3=1}^{s_3} \dots \sum_{j_n=1}^{s_n} u_1(e_1, e_{j_2}, e_{j_3}, \dots, e_{j_n}) \quad (4)$$

$$= \frac{s_1}{|S|} \sum_{\beta} u_1(e_1, e_{-\beta}) \quad (5)$$

which is exactly the average of the coefficients in the first row of A_1 . Thus for any agent i we have the equalities

$$\frac{s_i}{|S|} \sum_{\beta} u_i(e_1, e_{-\beta}) = \frac{s_i}{|S|} \sum_{\beta} u_i(e_2, e_{-\beta}) = \dots = \frac{s_i}{|S|} \sum_{\beta} u_i(e_{s_i}, e_{-\beta})$$

which after cancellation of the non-zero term $\frac{s_i}{|S|}$ proves our assertion. \square

Proposition 2. *If uniform distribution, $\hat{x} \in \Delta$, is a Nash equilibrium and in each of the payoff matrices the sums of the elements in each row that are larger than the average are equal, then it is an incentive equilibrium for $\varphi^D(x)$.*

Proof. We will use the above lemma to prove the assertion. Given that the rows must all have an equal sum, the average of the elements in A_i , which we will denote \bar{a}_i , is equal to $\frac{s_i}{|S|} \sum_{\beta} a_{i1\beta}$. Let us now consider the condition for an incentive equilibrium when our incentive is given by $\varphi^D(x)$. At a Nash

equilibrium we have the following calculation for each agent i

$$\varphi_{i\alpha}^D(\hat{x}) = \sum_{\gamma} (u_i(e_{i\alpha}, \hat{x}_{-i}) - u_i(\hat{x}) + a_{\alpha\gamma} - u_i(e_{i\alpha}, \hat{x}_{-i}))_+ \quad (6)$$

$$= \sum_{\gamma} (a_{\alpha\gamma} - u_i(e_{i\alpha}, \hat{x}_{-i}))_+ \quad (7)$$

$$= \sum_{\gamma} \left(a_{\alpha\gamma} - \frac{s_i}{|S|} \sum_{\beta} u_i(e_{i\alpha}, e_{-i\beta}) \right)_+ \quad (8)$$

$$= \sum_{\gamma} (a_{\alpha\gamma} - \bar{a}_i)_+ \quad (9)$$

where the second line is justified since \hat{x} is a Nash equilibrium and thus $u_i(e_{i\alpha}, \hat{x}_{-i}) = u_i(\hat{x})$. The last line is simply the sum of all the elements from row α that are larger than the average. Given our assumption, it must be the case that $\varphi_{i\alpha}^D(\hat{x}) = \varphi_{i\beta}^D(\hat{x})$ for every α and β . Thus we have $\varphi_{i\alpha}^D(\hat{x}) = \frac{1}{s_i} \sum_{\beta} \varphi_{i\beta}^D(\hat{x})$ for every agent i , which is true if and only if \hat{x} is an incentive equilibrium. \square

To summarize, we found a class of games where the Nash equilibrium coincides with the incentive equilibrium for $\varphi^D(x)$ at the uniform distribution. All RPS games have the property that the rows of the payoff matrices are permutations of the first row. Games with this property form a subset of the games where the Nash equilibrium and our incentive equilibrium agree.

We conjecture that this is the only agreement outside of constant games and strategies where players are receiving their respective maximum payoff. There are simple counterexamples when either of the conditions is

dropped. For example, if $A_i = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -3 & 1 \end{pmatrix}$, the average is 0, but the

sums across rows are not equal. The interior Nash equilibrium is $\hat{x} = ((1/6, 1/6, 2/3), (1/6, 1/6, 2/3))$ while the incentive equilibrium for φ^D is the uniform distribution. On the other hand, if $A_i = \begin{pmatrix} 1 & 2 \\ 3 & 0 \end{pmatrix}$ then the Nash equilibrium is the uniform distribution, but the incentive equilibrium is $\hat{x} \approx ((0.31, 0.69), (0.31, 0.69))$.

2 Asymptotic Stability

As we have seen, many of the dynamics that have Nash equilibria as fixed points do not necessarily converge to the uniform distribution. The specific examples that do (at least so far) have been Rock-Paper-Scissors type games. We notice that the main idea is to create a cycle of best replies by permuting the values in the first row of the payoff matrix. This cyclic behavior is essentially the problem with convergence. We will now show that changing the parameters while maintaining this type of cyclic payoff structure has no impact on the asymptotic stability of the incentive equilibrium for $\varphi^D(x)$.

Proposition 3. *If the rows of the payoff matrix A_i are permutations of each other, $\varphi_{i\alpha}^D(x) = \varphi_{i\beta}^D(x)$ for all $x \in \Delta$ and either $\varphi_{i\alpha}^D(\hat{x}) = 0$ or $\hat{x}_{i\alpha} = \frac{1}{s_i}$ for every α at incentive equilibrium.*

Proof. Denote σ as the permutation that takes row α to row β ; then every element in row β can be written as $a_{i\beta\gamma} = a_{i\alpha\sigma(k)}$ for some $k \in S_i$. Thus $\varphi_{i\beta}^D = \sum_{\gamma} (a_{i\beta\gamma} - u_i(x))_+ = \sum_k (a_{i\alpha\sigma(k)} - u_i(x))_+ = \varphi_{i\alpha}^D(x)$ regardless of x .

We can now use this fact to describe all possible incentive equilibria for $\varphi^D(x)$. By definition, at equilibrium \hat{x} , $\varphi_{i\alpha}^D(\hat{x}) = \hat{x}_{i\alpha} \sum_{\beta} \varphi_{i\beta}^D(\hat{x})$ for every i and every α . Given that the incentive functions are all equal regardless of $x \in \Delta$, we must have $\varphi_{i\alpha}^D(\hat{x}) = \hat{x}_{i\alpha} s_i \varphi_{i\alpha}^D(\hat{x})$, which is true if and only if $\varphi_{i\alpha}^D(\hat{x}) = 0$ for all α , which can occur only at the boundary or in a degenerate game, or when $\hat{x}_{i\alpha} = \frac{1}{s_i}$ for every α . \square

Recall the definition of an ISS is

$$\hat{x}_i \cdot \frac{\varphi_i^D(x)}{x_i} > x_i \cdot \frac{\varphi_i^D(x)}{x_i}$$

for all $x \neq \hat{x}$ in some neighborhood of \hat{x} . Also, an ISS is asymptotically stable wherever it satisfies the inequality in the definition. It will suffice then to prove that the uniform distribution is an ISS and the entire space is its basin of attraction.


Theorem 4. *If \hat{x} is a uniform incentive equilibrium for $\varphi^D(x)$ and the payoff matrices have rows that are permutations of each other, then \hat{x} is globally asymptotically stable in $\text{int}\Delta$ for the incentive dynamics.*

Proof. The previous proposition gives us that the incentives are equal for all α so we can without loss of generality use only $\varphi_{i1}^D(x)$ for each i .


$$\begin{aligned}
0 &> -\frac{1}{s_i} \sum_{\alpha} \frac{\varphi_{i\alpha}^D(x)}{x_{i\alpha}} + \sum_{\beta} \varphi_{i\beta}^D(x) \\
&= -\frac{\varphi_{i1}^D(x)}{s_i} \sum_{\alpha} \frac{1}{x_{i\alpha}} + s_i \varphi_{i1}^D(x) \\
&= \frac{\varphi_{i1}^D(x)}{s_i} \left[s_i^2 - \sum_{\alpha} \frac{1}{x_{i\alpha}} \right]
\end{aligned}$$

If we define $f(x) = \sum_{\alpha} \frac{1}{x_{i\alpha}}$ it is easy to show that $f(x)$ has a global minimum of s_i^2 when x_i is the uniform distribution. We simply optimize using Lagrange multipliers, noting that the Hessian matrix of $f(x)$ is positive definite in the interior of Δ . Thus \hat{x} satisfies the ISS definition for all $x \in \text{int}\Delta$. \square


We further conjecture that all interior incentive equilibrium are asymptotically stable. If this is true, we can reduce the open problem of finding a game dynamic where every orbit converges to a rest point to proving that the basins of attraction for the incentive equilibrium form a partition of Δ .



(a) BNN



(b) Logit (0.2)



(c) Smith

Figure 1: Stable limit cycles in Bad RPS




Figure 2: The replicator dynamics display invariant limit cycles and unstable equilibrium points in the RPS







Figure 3: The projection dynamics display invariant limit cycles and unstable equilibrium points in the RPS



(a) bad RPS



(b) RPS



(c) good RPS

Figure 4: Global asymptotic stability of the uniform distribution in the simultaneous updating dynamics

References

- [Fry12] D.E.A. Fryer, *On the existence of general equilibrium in finite games and general game dynamics*, Arxiv preprint arXiv:1201.2384 (2012).
- [SDF11] W.H. Sandholm, E. Dokumaci, and F. Franchetti, *Dynamo: Diagrams for evolutionary game dynamics, version 1.0*, <http://www.ssc.wisc.edu/~whs/dynamo>, 2011.